

SCIENCE on a GRAND SCALE

Through the Computing Grand Challenge Program, participants harness world-leading supercomputing and science to advance scientific discovery and Laboratory missions.

SINCE its inception in 2005, Lawrence Livermore's Computing Grand Challenge Program has awarded unclassified computing time on flagship institutional computing resources to scientists and engineers pursuing ambitious but achievable goals that advance national security or basic science research. (See *S&TR*, July/August 2011, pp. 13–19.) Now entering its second decade, the program is cultivating ideas and advances that are as important as ever. Fred Streitz, director of Livermore's High-Performance Computing Innovation Center and co-lead for the Grand Challenge Program, notes, "Grand Challenges are about encouraging people to think larger and giving them a venue to push the limits, scientifically and computationally."

These efforts have been aided by the ongoing evolution of high-performance computing (HPC). Clock speeds and numbers of processors have continued to climb, as exemplified by the changes in Grand Challenge resources, from the 20-teraflop (a teraflop equals 10^{12} floating-point operations per second), 4,000-processor Thunder supercomputer



Computational scientist Erik Draeger stands in front of a large, multiple-projector display of research simulating atomic-level behavior at the anode–electrolyte interface of a lithium-ion cell. Performed on Livermore’s Vulcan supercomputer, this simulation exemplifies the sort of cutting-edge research that is supported by the Computing Grand Challenge Program. (Photograph by George A. Kitrinos.)

used by Grand Challenge researchers in 2005 to the 5-petaflop (10^{15} flops), 400,000-processor Vulcan system available today. Architectures have also undergone more subtle changes, such as graphics processing units, in-processor memory, and custom central processing units. These trends have expanded the size and complexity of systems that can be modeled effectively. Innovations in computer architecture at Livermore have also expanded the scope of accessible research. For instance, this year's competition welcomed, for the first time, proposals for the 150-teraflop Catalyst supercomputer, which is especially designed to solve big-data analytics problems. "Catalyst is a different type of machine and allows Grand Challenges in a very different space than the others," says Streitz.

Although the computing resources are the enabler, the success of projects relies on the strength of the researchers' ideas. The program aims to inspire researchers not simply to scale up an existing solution, but to enhance or even reframe their investigations. Computer scientist Erik Draeger, a member of the Grand Challenge project selection committee, says, "The point of the program is to get scientists thinking in new directions. We have this

great hardware, but we're not making the most of it if we continue to think about problems in the same old way." Grand Challenge projects span many fields of research, including target simulations for the National Ignition Facility (NIF), which aid fusion energy science and stockpile stewardship research; first-principles molecular dynamics (MD) simulations, which support a wide range of basic and applied science efforts; and climate modeling, which has national and global security applications.

When Laser Meets Plasma

During fusion ignition experiments at NIF, multiple laser beams simultaneously enter through holes in a tiny metal container known as a hohlraum, striking the inside walls and producing x rays that compress the capsule of frozen fusion fuel in the hohlraum. To achieve the extreme temperatures and pressures such experiments require, laser light must be delivered to and absorbed by the hohlraum walls with great precision. Unfortunately, hohlraum conditions also invariably generate plasma, a gas of charged particles that can interfere with the experiment by misdirecting laser energy. Bent or scattered light can lower the temperature in the hohlraum and interfere


with symmetrical fuel compression and even cause facility damage if reflected back at the laser optics. (See *S&TR*, June 2010, pp. 17–19.)

The quest to understand, predict, and mitigate laser–plasma interactions (LPI) in NIF targets began 21 years ago, when NIF was little more than a collection of blueprints and plans, with the development of the three-dimensional (3D) LPI modeling code F3D, now called pF3D. Accurately simulating LPI is difficult and computationally intensive, as laser energy and plasma interact in complex, highly nonlinear ways. (See *S&TR*, January/February 2009, pp. 11–15; and January/February 2005, pp. 4–13.) The relevant phenomena span extremes in length and time, from macro to micro. Bridging these extremes is the computationally most challenging physics component—the mesoscale, the micrometer and picosecond scale behavior that pF3D seeks to understand.


Over the past two decades, LPI modeling has benefited from rapidly increasing computer performance at Lawrence Livermore. Bert Still, who served as pF3D's principal developer for many years, recalls a calculation he performed for NIF's groundbreaking ceremony in 1995: "The initial calculation was 128 by 128 by 512 cells and took an entire supercomputer to run. It only included forward light propagation and was pretty primitive by modern standards. Now I could do that calculation with the computing power in my cellular phone." At the time, researchers could only model a single thin beam filament near the entrance of the hohlraum interacting with homogeneous plasma.

Atlas Raises the Bar

LPI modeling efforts intensified in 2005, when a key external review of NIF technological readiness emphasized the need for 3D LPI simulations that could help evaluate potential ignition target designs. NIF was still under construction, and the laser power anticipated for ignition experiments was unprecedented. Pinpointing the target design and



Livermore's Computing Grand Challenge Program has seen growth in both requested and available time over its 10-year history, but demand has always exceeded supply.



The growth in computer performance at Livermore has enabled researchers to improve their laser-plasma interaction modeling capabilities. In 2009, for instance, researchers using an enhanced version of the code pF3D could perform calculations a million times larger—and therefore far more realistic—than those done in 1995 with an earlier version of the same code. The hohlraum illustrations on the bottom represent how greater supercomputing power has enabled increased complexity in simulations of National Ignition Facility (NIF) experiments.

experimental conditions most likely to meet the physics requirements while avoiding optics damage was imperative. Radiation hydrodynamics modeling expert Debbie Callahan observes, “We were dealing with a new machine, more energy, and a bigger plasma, all of which made prediction tough.” That year, using the Thunder supercomputer, the researchers were able to model the full two-millimeter-diameter laser beam, though still only for a short distance at the hohlraum entrance. Importantly, they simulated for the first time the large gradients in the plasma caused by temperature, electron density, and ion density variations that can affect the type and amount of light scatter that occurs.


In 2007 and 2008, Livermore LPI researchers received Grand Challenge awards totaling 16 million hours of machine time on the 44-teraflop Atlas supercomputer, enabling the team to take LPI simulation to a new level, both in size and amount of incorporated physics. Even with Atlas, the researchers could not simulate the whole beam, so they analyzed millimeter-scale radiation-hydrodynamics

simulations to determine when and where LPI was most likely to occur. Using the parameters set by these models and information gained from electron-scale modeling, they performed pF3D simulations of relevant spans of the laser beam’s path from the hohlraum entrance to its inner wall. Results were then fed back into the radiation-hydrodynamics model to help converge on a safe and optimized target design. These pF3D simulations provided motivation to lower the radiation temperature in the hohlraum, thereby reducing LPI effects, for instance. The most significant outcome of the Grand Challenge, however, was a modeling methodology for LPI prediction and target design evaluation that would prove vital as NIF experiments began in 2009.

LPI modeling expert Denise Hinkel says, “The Atlas simulations positioned us to be in the best space we could be when we turned on NIF. We were able to quickly address questions that arose.” They also laid the foundation for 2009 runs on the 500-teraflop Dawn supercomputer in which the researchers first simulated a full beam travelling the entire distance between the

hohlraum’s entrance and wall for roughly 100 picoseconds.

LPI modeling continues to shed light on phenomena not easily understood through experiments alone. The knowledge gained



The first-ever whole-beam pF3D simulation for a NIF ignition target was performed in 2009 on the Dawn supercomputer. Such simulations have helped ensure that target designs minimize misdirected light, which can impede the experiment and even damage the optics.


from simulations has, for instance, prompted NIF researchers to purposefully generate an LPI effect called crossbeam transfer and use it to achieve a more symmetrical implosion—a novel and effective approach. (See *S&TR*, March 2013, pp. 10–17.) According to Hinkel, the next true grand challenges in NIF modeling will be modeling LPI for multiple beams and integrating the three temporal and spatial scales of simulation, which will require computing at the exaflop (10^{18} flops) scale. The equivalent of 36 Dawns, for example, will be needed to model the interactions between plasma and several laser beams.

Take It from the Top

Standard MD simulations use empirical data in their calculations. (See *S&TR*, January/February 2006, pp. 11–17.) However, for complex systems or those under extreme conditions, gathering enough reliable experimental results to constrain all of the MD simulation's parameters can be difficult. An alternative approach for researchers is to bypass models and experiments and calculate the properties of materials from first principles—that is, directly from physics equations.

The Schrödinger equation, a multidimensional differential equation, can be used to understand the behavior of atoms and electrons at a quantum level, but solving the equation directly is computationally intractable for all but the smallest and simplest molecules, even on today's supercomputers. Instead, scientists rely on physical and numerical approximations for computing electronic structure, with the accuracy of these calculations depending on the approximations chosen. Even with approximations, the computational cost for first-principles MD is high, limiting the size and scope of what can be studied.

For instance, a decade ago, billion-atom classical MD calculations were common, while routine first-principle runs were restricted to just 50 atoms. Draeger notes, “A couple of hundred atoms was a huge challenge then. Now, with Grand Challenge



An image from a 25-picosecond, 1,700-atom Qbox simulation performed on the Vulcan supercomputer shows a lithium-ion cell anode–electrolyte interface. Qbox is a first-principles molecular dynamics (MD) code written to efficiently carry out large simulations on massively parallel supercomputers, allowing researchers to study systems with larger numbers of atoms than is possible with most other codes of its type.

resources, we can study thousands of atoms. Complexity in first-principles simulations goes up cubically, not linearly, so today we're solving problems that are thousands of times more difficult, on millions of times the resources.” The growth in HPC capabilities has enabled researchers to employ better approximations, study bigger problems, and explore new classes of chemical systems, including heterogeneous molecular interactions and extreme behavior such as shocks.

One of the codes that has allowed Livermore researchers to make strides in first-principles modeling is Qbox, an open-source application for which Draeger is lead Livermore developer. (See *S&TR*, November 2006, pp. 4–11.) Qbox uses the density functional theory (DFT) approximation approach, prized for its favorable ratio between precision

and computational cost. Most popular DFT codes were developed by and for academic researchers who only had access to hundreds or thousands of cores and so tend to run inefficiently on Livermore's massively parallel computing systems. Qbox, however, was designed for and thrives on such systems, allowing researchers to study larger numbers of atoms than is possible with most other DFT codes. Notes Draeger, “Qbox is a good example of how Grand Challenge support affects code design. It's a top-down rather than a bottom-up design.”

Putting Pressure on Hydrogen

Although DFT has proved a valuable and economical approach to first-principles MD simulations, some situations call for the greater accuracy permitted by quantum Monte Carlo (QMC). Unfortunately, QMC

is far more computationally expensive. (See *S&TR*, May 2007, pp. 4–11.) With support from the Grand Challenge and the Laboratory Directed Research and Development Programs, researcher Miguel Morales and his collaborators at the University of L'Aquila in Italy and the University of Illinois at Urbana-Champaign have been developing a more predictive first-principles approach that combines QMC and DFT. Their effort focuses on the behavior of hydrogen at extreme conditions—millions of degrees kelvin and millions of atmospheres. (See *S&TR*, January 2015, pp. 23–26.)

“We chose hydrogen due to its high impact and because it’s a simple enough element that we can bring all of the techniques we’ve developed over the last decade to bear, in terms of describing properties from computer simulation without experimental input. It’s still a very ambitious project that requires the largest computer resources we can get,” says Morales. Thanks to the Grand Challenge Program, those resources have included the 260-teraflop Sierra and 5-petaflop Vulcan systems. The project not only refines an important predictive computational method but also aims to shed light on planetary formation. Gas giants, such as Jupiter, are over 90 percent hydrogen and helium, and pressures and temperatures within the planets can vary by orders of magnitude, necessitating an accurate phase diagram covering a large range of thermodynamic conditions.

Technological limitations have also restricted the experimental data available for high-pressure hydrogen, a regime where DFT first-principles methods have struggled quantitatively. “To correctly model the interior of Jupiter, we need to know when and how certain transitions happen,” explains Morales, “but that is precisely where DFT becomes inaccurate.” Morales’ solution has been to use a combination of methods. QMC is used to check the accuracy of DFT calculations, particularly around key transitions, such as when dense liquid


hydrogen changes from metal to insulator. “The more expensive method acts as the decision-making guide for the less expensive one,” he adds.

This approach is enabling scientists to more precisely pinpoint when and how phase transitions occur, including the metallization of solid hydrogen, a high-temperature superconductor candidate. “At this point,” says Morales, “we have results that match very well with our experimental data. Such data only exists in small slivers, so we benchmark our methods on the experiments and then go and explore areas that haven’t been explored in experiments.” While Morales has been successfully demonstrating a hybrid approach to first-principles MD

simulations, other Livermore researchers have begun examining how to reduce the resource intensiveness of the calculations themselves—an ambitious effort. “The next major effort at Livermore will be to create new algorithms that knock the complexity down from n^3 to n , like classic MD,” says Draeger.

New Heights in Climate Modeling

Climate modeling has benefited greatly from the steep climb in computing capability over the past several decades. In 1998, a 1-year simulation of global climate using 300-kilometer resolution could be run in a day on a supercomputer. The same model can now be run on a high-end



Livermore researchers are using a combination of first-principles MD approaches to better understand the phase diagram of hydrogen, particularly at high pressure, where it exhibits transitions such as from a molecular solid to a quantum liquid, as this artist’s rendering suggests. This research will likely shed light on the formation of planets such as Jupiter and Saturn, which are mostly composed of hydrogen and helium.

desktop machine in minutes. Climate scientists have responded by creating more detailed and accurate simulations, using finer resolutions, more variables, and longer time spans. The additional computing horsepower is also allowing them to refine methods for characterizing and reducing uncertainty, giving them greater confidence in their modeling projections. At Livermore, a series of Grand Challenge efforts has accelerated the development of these higher performance models and uncertainty quantification methods, with which climate researchers can gain new insights into Earth's climatic past, present, and future.

The first of these Grand Challenge projects, performed from 2005 to 2007 on the Thunder supercomputer, evaluated a computationally intensive approach for improving regional climate prediction called dynamical downscaling. This approach uses high-resolution simulations of global climate to drive even higher resolution regional-scale models. Govindasamy Bala's team employed the primary U.S. general circulation model, known as the Community Climate System Model (CCSM), to perform a 400-year global simulation with improved results for global surface winds and sea surface temperatures. At 100-kilometer resolution, this was the highest resolution multicentury CCSM simulation performed to date.

Results of the global simulation served as initial conditions and boundary data for

12-kilometer-scale calculations of climate change in California over four decades using the Weather and Research Forecasting model. Both the resolution and duration of the regional simulations were unparalleled in regional climate modeling. The resolution enabled the team to resolve more of California's complex topography, such as details in mountainous areas, and the length of the run allowed for a better sampling of natural variability, increasing the credibility of model predictions. Comparison with observational datasets and global modeling data confirmed that dynamical downscaling provides valuable insights into regional climate behavior, such as small-scale atmospheric features that cannot readily be captured in a global model.

Climate Clarity through Uncertainty


CCSM also featured heavily in a 2008 Grand Challenge effort led by Dave Bader. Using 8 million hours of Atlas processor time to complete a simulation of the global climate under present-day conditions—a necessary precursor to climate projection—the team achieved what Bader describes as “an unprecedented realism of phenomena.” The 20-year simulation was configured using grid resolutions of 11 kilometers for the ocean and sea ice and 28 kilometers for the atmosphere and land. This was the first study with fine enough oceanic and atmospheric horizontal resolution to simulate turbulent instabilities in the large-scale circulation—for instance, the

formation and propagation of tropical cyclones, the frequency and intensity of which many researchers project will be affected by climate change. The major outcome of the project, though, was a process for performing ultrahigh-resolution global climate modeling. “Ten years ago, the big challenge was weather-scale climate modeling,” says Bader. “The Grand Challenge project put us on the path to doing this kind of modeling routinely.”


A third Grand Challenge, led by Richard Klein from 2009 to 2011, applied rigorous and computationally demanding uncertainty quantification methods developed for the Stockpile Stewardship Program to climate prediction. (See *S&TR*, July/August 2010, pp. 12–14.) Uncertainty can take many forms. Simply using climate models to assess climate change introduces uncertainty because the models do not perfectly represent the climate system, and the various models respond differently to the same input. In fact, more than 100 parameters, each with associated uncertainties, can influence climate simulation predictions. Uncertainty quantification is performed with an ensemble of models. Klein's team is running CCSM on Atlas and Sierra—along with an intelligent, self-adapting search tool they developed—to generate a comprehensive set of climate simulations and comb through possible combinations of input parameters. This methodology is enabling researchers to pinpoint, measure, and potentially reduce

The plot shows modeled (middle and right) and observed (left) surface wind circulation in the summertime in the Arctic. Using finite-volume transport, an improved method for simulating global surface winds and sea surface temperatures, Livermore global climate simulations were able to better simulate features such as cyclonic circulation than with traditional modeling methods.


Analysis of observed data



Spectral simulation



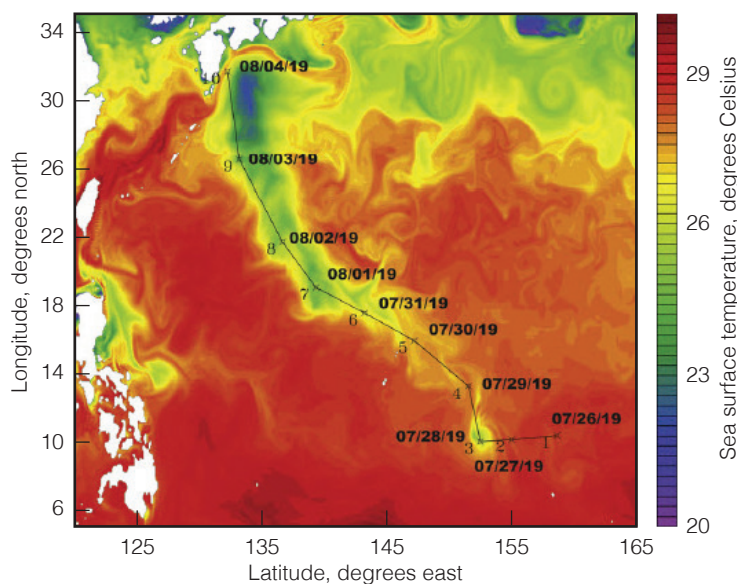
Livermore finite-volume simulation



sources of prediction uncertainty, as well as to assess low-probability but high-consequence events such as rapid melting of the polar ice sheets.

Livermore Grand Challenge simulations and calculations helped lay the foundation for new climate research efforts, most notably the Accelerated Climate Modeling for Energy initiative, launched in 2014 by the Department of Energy's Office of Science. Over the next decade, the initiative's academic, industry, and government partners, including Lawrence Livermore, will expedite the development and testing of models for climate and energy applications. This work is done in anticipation of new and disruptive HPC architectures, such as those to be delivered through the Collaboration of Oak Ridge, Argonne, and Livermore in 2017 and 2018 and the exaflop-scale machines likely to follow. (See *S&TR*, March 2015, pp. 11–15.) These models will have the same resolution as Bader's Grand Challenge simulation but will be more sophisticated and computationally demanding Earth system models. "Earth system models are the future," says Bala. "These are models that include carbon, nitrogen, sulfur, and phosphorous cycles besides the usual physical components."

States Bader, "The big challenges for climate modeling are actually harder now than those we faced a decade ago. Climate scientists have shown that climate change is real, increasing, and potentially irreversible. Now policymakers are asking for tools to predict the rate of change and answer other hard questions." Climate scientist Ben Santer adds, "We know that Earth's climate system is going to experience profound changes, such as large-scale warming and moistening of the atmosphere, rising sea levels, retreat of snow and sea-ice cover, and increases in the frequency and intensity of heat waves, but the regional and seasonal details of these changes are much fuzzier." Predicting these details with precision and confidence and delivering information that can help countries and communities



This ultrahigh-resolution global simulation of present-day climate conditions was detailed enough to track the progress of a Category 4 hurricane and its associated wake of cold water in the tropical northwest Pacific. Modeling extreme climate events is an important capability, as their frequency and intensity is expected to be increased by climate change.

make resource-planning decisions will require enhanced models and exaflop-scale computing capabilities.

Building on a Decade of Success

By most any measure, the Grand Challenge Program's inaugural decade has been a success. The program has seen steady growth in proposals and time requested, and the quality of ideas evaluated each year by internal and external referees has been consistently high. Furthermore, many boundary-pushing concepts have gone on to become new, robust programs or projects or to boost existing ones. Multiprogrammatic and Institutional Computing Program director and Grand Challenge Program co-lead Brian Carnes notes, "When a project is granted significant HPC resources to fully develop its science and technology, that's when it starts having value. The Grand Challenge Program allows that value to develop, so that it can be impactful to the programs." Pacific Northwest, Los Alamos, and Sandia national laboratories have also created their own Grand Challenge-type programs modeled on Lawrence Livermore's.

The Grand Challenge Program represents a sizeable investment of the Laboratory's

computing time. In fact, only six countries in the world possess more computing resources than Livermore makes available to individual researchers and programs for unclassified computing through the Multiprogrammatic and Institutional Computing Program initiatives, including the Computing Grand Challenge Program. These investments pay dividends by advancing challenging and mission-relevant research and honing the skills of researchers through access to new computational architectures and modeling and simulation methods.

—Rose Hansen

Key Words: Accelerated Climate Modeling for Energy, Atlas, climate model, Community Climate System Model (CCSM), Computing Grand Challenge Program, Dawn, density functional theory (DFT), Earth system model (ESM), exaflop, first-principles model, high-performance computing (HPC), hohlraum, laser-plasma interaction (LPI), molecular dynamics (MD), Multiprogrammatic and Institutional Computing Program, National Ignition Facility (NIF), petaflop, Qbox, quantum Monte Carlo (QMC), Sierra, teraflop, Thunder, uncertainty quantification, unclassified computing, Vulcan.

For further information contact Brian Carnes (925) 423-9181 (carnes1@llnl.gov).